

WHAT IS CLAIMED IS:

1 1. A data processing node configured to participate in global transactions
2 comprising:

3 A) a transaction processing application operating within a global transaction
4 processing protocol incorporating a two-phase commit procedure;

5 B) coupling means to at least one other data processing node operating
6 within the global transaction processing protocol for information exchange
7 therewith;

8 C) a memory storing database files subject to access during transaction
9 processing;

10 D) a database access application communicating with the transaction
11 processing application to address files stored in the memory and transfer
12 data between the transaction processing application and the memory;

13 E) a file management system having exclusive access to reserved locations
14 in the memory for reading and writing meta-data therein;

15 F) physical file access logic selectively coupling the memory and the
16 database access application, the physical file access logic incorporating file
17 protections which are controlled by the file management system; and

18 G) node recovery means for recovering following a system failure which
19 occurs during transaction processing on a database stored in said memory

and faithfully rebuilding, after restart, the data processing node to the state immediately prior to the failure; said node recovery means including:

- 1) means to remove the results of incomplete non-“in-doubt” transactions;
- 2) means to lock files which have been updated by “in-doubt” transactions; and
- 3) means for permitting normal access to the database after the operations carried out by the means set forth in subparagraphs G)1) and G)2) have completed.

2. The data processing node of Claim 1 in which the meta-data for each file stored in memory during transaction processing includes an undo log, and the undo log entry for each file affected by a data processing node failure is accessed after restart to provide information essential to rebuilding the state of the transaction at the data processing node at the time of the failure.

3. The data processing node of Claim 1 in which the rebuild effected is to the “in-doubt” state.

1 4. The data processing node of Claim 2 in which the rebuild effected is to the "in-
2 doubt" state.

1 5. The data processing node of Claim 2 in which "recovery-needed", "recovery-
2 done" and "in-doubt" counts are normally maintained in individual file description
3 records and in which, in the event of a data processing node failure as evidenced
4 by a difference between the BSI of the current system boot and the BSI as
5 recorded in the file description record:

6 A) if the node failure occurs when there is a non-zero "in-doubt" count,
7 then, during restart:

8 1) if the "recovery-needed" is non-zero, both the "recovery-done" and
9 "in-doubt" counts are cleared; and

10 2) if the "recovery-needed" is zero and the "in-doubt" is non-zero, the
11 original "recovery-needed" count is reinstated.

1 6. The data processing node of Claim 5 in which the rebuild effected is to the "in-
2 doubt" state.